SYMBOL GROUNDING ASSOCIATION IN MULTIMODAL SEQUENCES WITH MISSING ELEMENTS

Federico Raue[1,2], Thomas M. Breuel[1], Andreas Dengel[1,2], Marcus Liwicki[1]
[1]University of Kaiserslautern, Germany
[2] German Research Center for Artificial Intelligence (DFKI), Germany.
{federico.raue,andreas.dengel}@dfki.de, {tmb,liwicki}@cs.uni-kl.de

arXiv:1511.04401v3 [cs.CV] 12 Jan 2016

ABSTRACT

In this paper, we extend a symbolic association framework to being able to handle missing elements in multimodal sequences. The general scope of the work is the symbolic associations of object-word mappings as it happens in language development on infants. This scenario has been long interested by Artificial Intelligence, Psychology and Neuroscience. In this work, we extend a recent approach for multimodal sequences (visual and audio) to also cope with missing elements in one or both modalities. Our approach uses two parallel Long Short-Term Memory (LSTM) networks with a learning rule based on EM-algorithm. It aligns both LSTM outputs via Dynamic Time Warping (DTW). We propose to include an extra step for the combination with max and mean operations for handling missing elements in the sequences. The intuition behind is that the combination acts as a condition selector for choosing the best representation from both LSTMs. We evaluated the proposed extension in three different scenarios: audio sequences with missing elements, visual sequences with missing elements, and sequences with missing elements in both modalities. The performance of our extension reaches better results than the original model and similar results to a unique LSTM trained in one modality, i.e., where the learning problem is less difficult.

1 INTRODUCTION

A striking feature of the human brain is to associate abstract concepts with the sensory input signals, such as visual, audio. As a result of this multimodal association, a concept can be located and translated from a representation of one modality (visual) to another representation of a different modality (audio), and vice versa. For example, the abstract concept "ball" from the sentence "John plays with a ball" can be associated to several instances of different spherical shapes (visual input) and sound waves (audio input). Several fields, such as Neuroscience, Psychology, and Artificial Intelligence, are interested to determine all factors that are involve in this task. However, this association is still an open problem (Steels, 2008). This scenario is known as *Symbol Grounding Problem* (Harnad, 1990).

With this in mind, infants start learning the association when they are acquiring the language(s) in a multimodal scenario. Gershkoff-Stowe & Smith (2004) found the initial set of words in infants is mainly nouns, such as dad, mom, cat, and dog. In contrast, the language development can be limited by the lack of stimulus (Andersen et al., 1993; Spencer, 2000), i.e. deafness, blindness. Asano et al. (2015) found two different patterns in the brain activity of infants depending on the semantic correctness between a visual and an audio stimulus. In simpler terms, the brain activity is pattern 'A' if the visual and audio signals represent the same semantic concept. Otherwise, the pattern is 'B'.

Related work has been proposed in different multimodal scenarios inspired by the Symbol Grounding Problem. Yu & Ballard (2004) explored a framework that learns the association between objects and their spoken names in day-to-day tasks. Nakamura et al. (2011) introduced a multimodal categorization applied to robotics. Their framework exploited the relation of concepts in different modalities (visual, audio and haptic) using a Multimodal latent Dirichlet allocation. Previous ap-

proaches has focused on feature engineering, and the segmentation and the classification tasks are considered as independent modules.

This paper is focused on an end-to-end system in order to segment and classify the object-word association problem. Moreover, we are interested in multimodal sequences where the semantic concept sequence in one modality has missing elements in the other modality. For instance, one modality sequence is represented by "acf", and the other modality is represented by "cdf". Also, the association problem is different from the traditional setup where the association is fixed via a pre-defined coding scheme of the classes (e.g. 1-of-K scheme) before training. We explain the difference of our association problem setup and common approaches in multimodal machine learning in Section 1.1. In this work, we are evaluation the advantage of max and mean operation for handling missing elements in one or two modalities. Our work is an extension from Raue et al. (2015) where both modalities represent the same semantic sequence (no missing elements). In short, the model was implemented by two Long Short-Term Memory (LSTM) networks that their outputs were aligned using Dynamic Time Warping.

This paper is organized as follows. We shortly describe Long Short-Term Memory networks, mainly sequence classification in unsegmented inputs, in Section 2. The original end-to-end model for object-word association is explained in Section 3. The max and mean steps are presented in Section 4. A generated dataset of multimodal sequences with missing elements is described in Section 5. In Section 6, we compare the performance of the proposed extension against the original model and a single LSTM network trained on one modality in the traditional setup.

1.1 Multimodal Machine Learning

Recent results in Machine Learning show several successful applications in multimodal scenarios. We want to indicate the difference between previous tasks and the presented task.
Feature Fusion: Deep Boltzmann Machines have been applied to generate multimodal features in unsupervised environment. (Srivastava & Salakhutdinov, 2012; Sohn et al., 2014).
Textual Description: Convolutional Neural Networks (CNN) in combination with LSTM have been already applied in cross-modality scenario, which *translates* images to textual captions (Vinyals et al., 2014; Karpathy et al., 2014).
Object-Word Association: In this work, we define the association between objects and words without specifying a coding scheme for each class before training. With this in mind, two definitions are introduced for explanation purposes: semantic concepts (SeC) and symbolic features (SyF). We explain these definitions through the following example. A classification problem is usually defined by an input that is associated with a class (or a semantic concept) and this class is represented by a vector (or a symbolic feature). That relation is fixed and it is chosen externally (outside of the network). In contrast, the presented task includes the relation between the class and its corresponding vector representation as a learned parameter. At the end, the model not only learns to associate objects and words but also learns the *symbolic structure* between the semantic concepts and the symbolic features. Figure 1 shows examples of each component and the differences between the traditional setup and our task. It can be seen that our task involves only three components, whereas the traditional setup uses only two components (red box).

2 Background: Long Short-Term Memory (LSTM) networks

Long Short-Term Memory (LSTM) is a recurrent neural network proposed by (Hochreiter & Schmidhuber, 1997). LSTM is capable to learn long sequences without the vanishing gradient problem (Hochreiter, 1998). This is accomplished based on memory cells and gates that learn when to read, keep or reset. Graves et al. (2006) introduced Connectionist Temporal Classification (CTC) for learning sequences in one-dimension from unsegmented inputs. They successfully applied it to speech recognition. Their idea was to add an extra element (*blank class (b)*) to the original sequence in order to align two sequences. One sequence was the output from LSTM and the other sequence is obtained by forward-backward[1] algorithm to LSTM output. A decoding mechanism was required in order to convert LSTM output to classes. More information can be found in the original paper (Graves et al., 2006).

[1] similar for training Hidden Markov Model

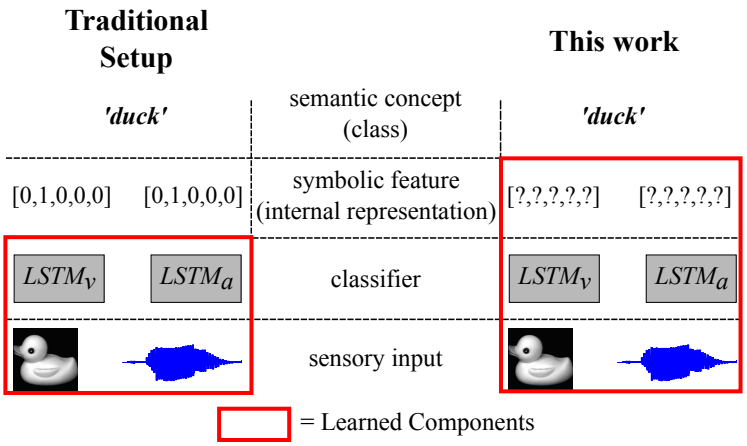

Figure 1: Comparison of components between the traditional setup and our setup for associating two multimodal. Note that our task has an extra learnable component (relation between semantic concepts and their representations), whereas the traditional scenario is already pre-defined (red box).

3 Multimodal Symbolic Association

In Section 1, we mentioned that this work is based on Raue et al. (2015). They introduced the symbolic association in multimodal sequences. Moreover, their initial assumption was a one-to-one relation between the representation of the same semantic concepts in two modalities, i.e., visual and audio. For example, the semantic concept sequence '1 2 3' is visually represented by a text line of digits and auditory represented by the waveforms of the spoken words (one two three). Furthermore, the model was implemented using two parallel LSTM networks and an EM-style training rule. In more detail, the model exploited the recent results of LSTM in segmentation and classification tasks, and the EM training rule was applied for learning the agreement under the proposed constraints.

In summary, the training rule worked as follows. Initially, each modality was fed to each LSTM network. Then, a statistical constraint was applied to the output of LSTM for determining the symbolic structure. Here, the symbolic structure was the relation between semantic concepts and symbolic features. For instance, the semantic concept "duck" can be represented by the index "4" (1-of-K coding scheme). Afterwards, the output of LSTM and the found symbolic structure were combined for feeding to a forward-backward algorithm (one component in CTC layer in Section 2). At this point, both LSTM networks worked independently in each modality. For association both networks, the outputs from both forward-backward steps are aligned using Dynamic Time Warping (DTW). Therefore, the information of one modality is translated to the other modality, and vice versa. It is important to realize that one LSTM used the other LSTM as a target for updating the weights. Figure 2 illustrates the training algorithm with an example.

We shortly explain the statistical constraint and DTW modules in Sections 3.1 and 3.2. Please refer to the original paper for more details (Raue et al., 2015).

3.1 Statistical Constraint

The goal of this module was to predict the symbolic structure between the semantic concepts and the symbolic features. With this in mind, a set of concept weights were defined ($\boldsymbol{\gamma_1}, \ldots, \boldsymbol{\gamma_c}; c \in SeC$), where $\boldsymbol{\gamma_1}, \ldots, \boldsymbol{\gamma_c}$ were vectors with $|SeC|$ elements, and SeC was the set of all semantic concepts. In addition, this learnable component was trained in an EM-style algorithm.

First, *E-step* predicted the symbolic structure (Γ^*) given the output of LSTM and the concept weights. This was defined by

$$\boldsymbol{\Gamma_c} = \frac{1}{T} \sum_{t=1}^{T} f\left(\boldsymbol{z_t}, \boldsymbol{\gamma_c}\right) \tag{1}$$

$$\boldsymbol{\Gamma^*} = g\left(\left[\boldsymbol{\Gamma_1}, \ldots, \boldsymbol{\Gamma_c}\right]\right) \tag{2}$$

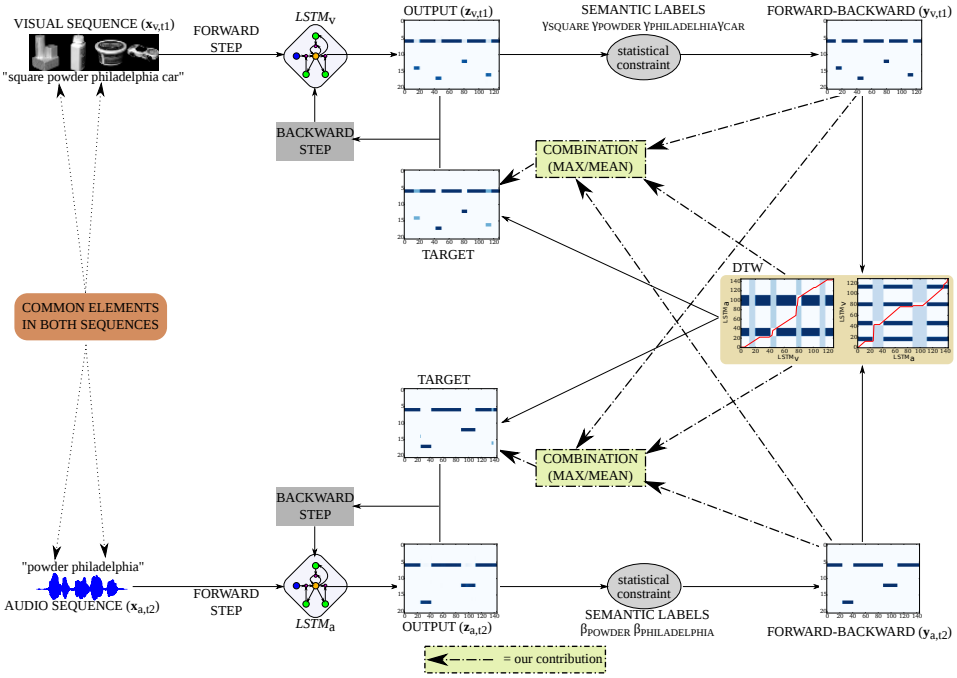

Figure 2: General overview of the original model and our contribution. In this work, we include a module (max/mean operation) that combines the forward-backward step from one LSTM and the aligned forward-backward step from the other LSTM.

where $\boldsymbol{z_t}$ was the output vector of LSTM with $|SeC|$ elements at time t, $\boldsymbol{\gamma_c}$ was the vector of concept weight c, *function* f was the element-wise power operation between $\boldsymbol{z_t}$ and $\boldsymbol{\gamma_c}$. *Function* g was a max row-column elimination, where the maximum element $\Gamma_{i,j}$ was set to one, and the rest of values at column j and row i were set to zero. This is repeated $|SeC|$ times.

Second, *M-step* updated the concept weights given the output of LSTM and the symbolic structure. As a result, the following cost function is defined

$$\boldsymbol{cost_c} = \left(\frac{1}{T} \sum_{t=1}^{T} f\left(\boldsymbol{z_t}, \boldsymbol{\gamma_c}\right) - \frac{1}{|SeC|} \boldsymbol{\Gamma_c^*} \right)^2 \tag{3}$$

$$\boldsymbol{\gamma_c} = \boldsymbol{\gamma_c} - \alpha * \nabla_c \, \boldsymbol{cost_c} \tag{4}$$

where $\boldsymbol{\Gamma_c^*}$ is the column vector of $\boldsymbol{\Gamma^*}$ that represents the relation between the semantic concept c and its corresponding symbolic feature, α was the learning rate, and $\nabla_c \, \boldsymbol{cost_c}$ was the derivative w.r.t. to $\boldsymbol{\gamma_c}$.

3.2 Dynamic Time Warping (DTW)

This module exploited the monotonic behavior of the multimodal sequences in one dimension. In that case, the output of each forward-backward module was aligned using Dynamic Time Warping (DTW) (Berndt & Clifford, 1994). The matrix DTW is calculated based on the following path

$$DTW[i,j] = dist[i,j] + min \begin{cases} DTW[i-1, j-1] \\ DTW[i-1, j] \\ DTW[i, j-1] \end{cases} \tag{5}$$

where $dist[i,j]$ was the euclidean distance between the outputs at timestep i of one LSTM and at timestep j of the other LSTM. As a result, the forward-backward algorithm from one network can be translated to the other LSTM based on this alignment. In other words, the loss functions for updating the LSTM weights were defined by

$$loss_function_v = y_{a,t_2}[aligned_to_LSTM_v] - z_{v,t_1} \tag{6}$$

$$loss_function_a = y_{v,t_1}[aligned_to_LSTM_a] - z_{a,t_2} \quad (7)$$

where z_{v,t_1} and z_{a,t_2} were the output of LSTM (visual and audio, respectively) at timestep t_1 and t_2 (respectively), y_{v,t_1} and y_{a,t_2} were the output of the forward-backward algorithm of each LSTM, and $aligned_to_LSTM_v$ and $aligned_to_LSTM_a$ were the paths from DTW cost matrix.

4 MAX AND MEAN OPERATIONS

In this paper, we are interested in the symbol grounding in multimodal sequences contains missing elements in one or two modalities. A formal definition of problem is giving by $\mathcal{M} = \{(\boldsymbol{x_v}, \boldsymbol{p_{1,...,n}}), (\boldsymbol{x_a}, \boldsymbol{q_{1,...,m}})\}$ where $\boldsymbol{x_v}$ and $\boldsymbol{x_a}$ are sequences of visual and audio elements, $p_{1,...,n}$ and $q_{1,...,m}$ are the correspondence semantic concept sequence of length n and m. Each modality can have elements that are or are not presented in the other modality. We want to indicate this scenario is different from the previous work, whereas both multimodal sequences have the same semantic concepts.

On one hand, $LSTM_v$ receives the visual sequence $\boldsymbol{x}_v$ and the visual semantic sequence $p_{1,...,n}$. On the other hand, $LSTM_a$ receives the audio sequence $\boldsymbol{x}_a$ and the audio semantic sequence $q_{1,...,m}$. Then, each LSTM output and their corresponding semantic sequences are fed to the statistical constraint module. Afterwards, each symbolic structure is used for applying the forward-backward algorithm. In the original paper, the *loss functions* is calculated by Equations 6 and 7. We are proposing to combine the output of one network and the aligned output from the other network. As a result, the *loss functions* are updated:

$$loss_function_v = h(y_{a,t_2}[aligned_to_LSTM_v], y_{v,t_1}) - z_{v,t_1} \quad (8)$$

$$loss_function_a = h(y_{v,t_1}[aligned_to_LSTM_a], y_{a,t_2}) - z_{a,t_2} \quad (9)$$

where *function h* can be the max or mean element-wise operation. The intuition behind is to give the common elements in the multimodal sequence the best probabilities regardless of the modality and to keep the best probability of the current modality when the elements are missing.

5 EXPERIMENTAL DESIGN

5.1 DATASETS

We generated three different multimodal datasets where the elements of the sequence are missing in one or both modalities, but the relative order between the elements is the same. For example, a visual semantic concept sequence can be represented by a text line of digits "2 4 7", and an audio semantic concept sequence can be represented by "two seven". In this case, we assumed a simplified scenario of symbol grounding, where the continuity of semantic concepts is different on each modality. The visual component is a horizontal arrange of isolated objects, and the audio component is a speech that represents some elements of the visual component, and vice versa. We want to point out that the visual component is similar to a panorama view. The procedure for generating the multimodal datasets is explained.

Visual Component: We used COIL-20 (Nene et al., 1996) that is a standard dataset of 20 isolated objects with a black background. Each isolated object has 72 views at different angles. After selecting the object for the sequence, each object is rescale to 32 x 32 pixels and stacked horizontally. The odd angles were used for training, and the even angles were used for testing.

Audio Component: We generated wav files using Festival Toolkit (Taylor et al., 1998) given the semantic concepts in the sequence. Additionally, we used five artificially voices in English that added to Festival.

Generating Semantic Sequences: Two scenarios were considered for generating the semantic sequences for each modality: missing elements in one modality and both modalities. For the first scenario, we generated randomly sequences between 4 and 8 semantic concepts. Later, the missing elements were randomly selected between 1 and maximum 50% of the original length of the sequence. For the second scenario, we changed the rule to randomly selecting elements between 1 and 25% of the original length in each modality. Thus, the multimodal sequences were different. As a

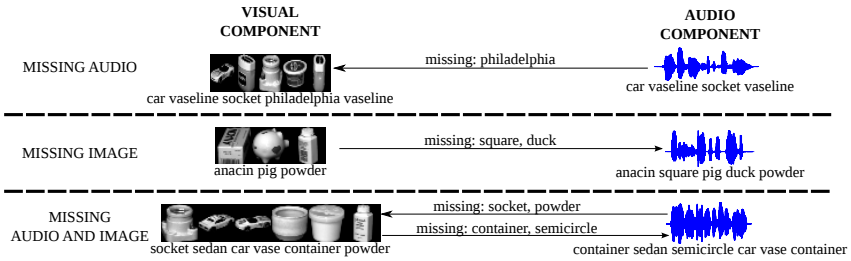

Figure 3: Samples of the three multimodal datasets.

Table 1: Label Error Rate (%) from the three multimodal datasets. It can be seen that the original model performs worse than the proposed combination in the three scenarios. Also, the presented extension reached similar results to LSTM.

DATASET	METHOD	VISUAL COMPONENT	AUDIO COMPONENT
MISSING ELEMENTS IN VISUAL COMPONENT	LSTM	0.21 ± 0.16	0.28 ± 0.11
	Original Model	8.47 ± 2.02	13.88 ± 2.14
	Original Model + MAX	0.67 ± 1.06	0.75 ± 1.58
	Original Model + MEAN	0.73 ± 0.65	0.82 ± 1.12
MISSING ELEMENTS IN AUDIO COMPONENT	LSTM	0.25 ± 0.22	0.28 ± 0.12
	Original Model	5.88 ± 1.51	1.04 ± 0.52
	Original Model + MAX	0.65 ± 0.30	0.09 ± 0.03
	Original Model + MEAN	0.92 ± 0.38	0.31 ± 0.38
MISSING ELEMENTS IN BOTH MODALITIES	LSTM	0.20 ± 0.10	0.28 ± 0.08
	Original Model	16.93 ± 14.05	14.72 ± 5.62
	Original Model + MAX	0.57 ± 0.57	0.23 ± 0.09
	Original Model + MEAN	0.64 ± 0.36	0.21 ± 0.06

result, each modality is compounded by common and missing elements. We chose the following vocabulary for each semantic concepts: *duck, triangle, sedan, cat, anacin, car, square, powder, tylenol, vaseline, semicircle, glass, pig, socket, container, bottle, vase, cup, convertible, philadelphia*.

Training and Testing Datasets: We generated a multimodal dataset of 50.000 sequences for training and 30.000 sequences for testing. Some examples of the three multimodal datasets are shown in Figure 3.

5.2 Input Features and LSTM setup

We did not apply any pre-process step for the visual component. In contrast, the audio component was converted to Mel-Frequency Cepstral Coefficient (MFCC) using HTK toolkit[2]. The audio representation was a vector of 123 components: a Fourier filter-bank with 40 coefficients (plus energy), including the first and second derivatives. Afterwards, all audio components were normalized to have zero mean and standard deviation one.

The proposed extension was compared against the original model. Also, we compared the extension against LSTM with CTC layer. The parameters of the visual LSTM were: 20 memory cells, learning rate 1e-5, and momentum 0.9. On the other hand, the audio LSTM had 100 memory cells, and the learning rate and momentum are the same as in the visual LSTM. In addition, the learning rate in the statistical constraint was set to 0.001.

6 Results and Discussion

As mentioned previously, the assumption of the original model was to represent the same semantic concept sequence in both modalities. In other words, a one-to-one relation existed between modali-

[2]http://htk.eng.cam.ac.uk

ties. In contrast, our assumption is more challenging because the semantic concept in one modality can be or cannot be in the other modality. We reported the Label Error Rate (LER) as a performance metric, which is defined by

$$LER = \frac{1}{|Z|} \sum_{(x,y) \in Z} \frac{ED(x,y)}{|y|} \tag{10}$$

where x is the output classification, y is the ground-truth, and $ED(x,y)$ is the edit distance between the output classification and the ground-truth. In addition, LSTM was trained using 10.000 random samples from the training set and tested using 1.000 random samples from the testing set. This procedure was applied ten times.

Table 1 summarizes the performance of LSTM, the original model, and the presented extension. Those results are divided into two parts as follows. First, the proposed extension handles missing elements in multimodal sequences better than the original model. It can be infer that the max and mean operations keep the strongest representation between modalities. Note that, the representations are used for updating the weights in the backward step. Second, the proposed extension reaches similar results to the standard LSTM. In this case, LSTM was trained in each modality independently. As a reminder, we mentioned two setups for classification tasks: the traditional setup and the setup used in this work. In other words, the proposed extension decreased the performance a little because of the extra learned parameter (included in our constraint). In summary, the max operator reached the best results in the three multimodal datasets.

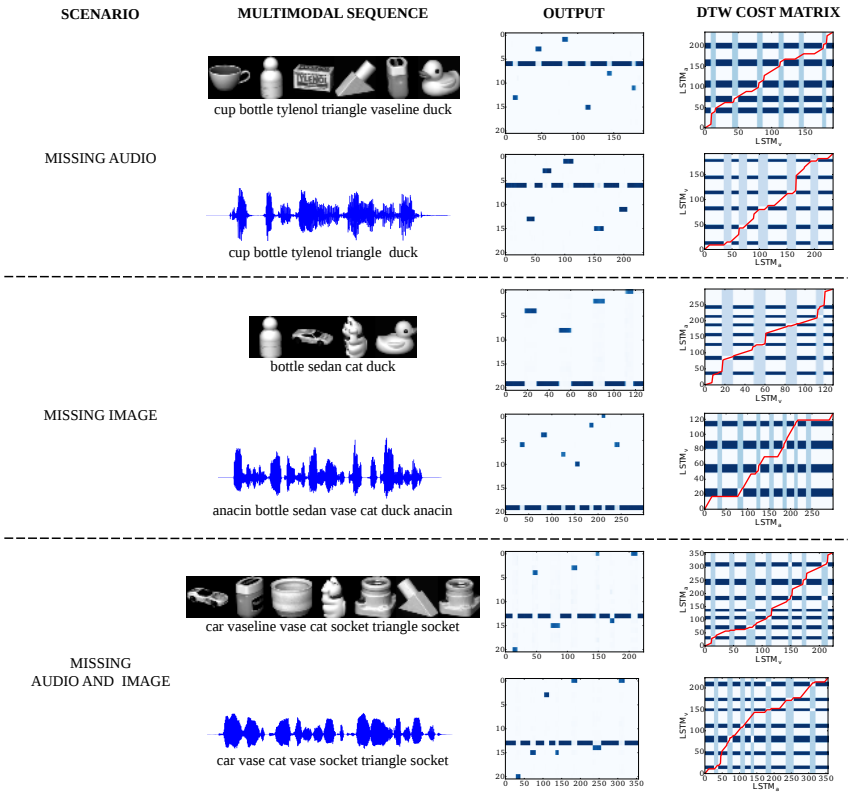

Figure 4: Example of several LSTM outputs and their corresponding cost matrix. Note that, the common semantic concepts have the right position and the same classification in both LSTMs.

Another outcome in this work is the agreement of the symbolic structure in both modalities, even with missing elements. Figure 4 shows examples of the agreement in three different cases. It can be observed that both LSTMs learn to segment and classify the object-word relation in unsegmented multimodal sequences. Moreover, the common concepts in both modalities are represented by a similar symbolic feature and located at the right position on the sequence. For example, the semantic

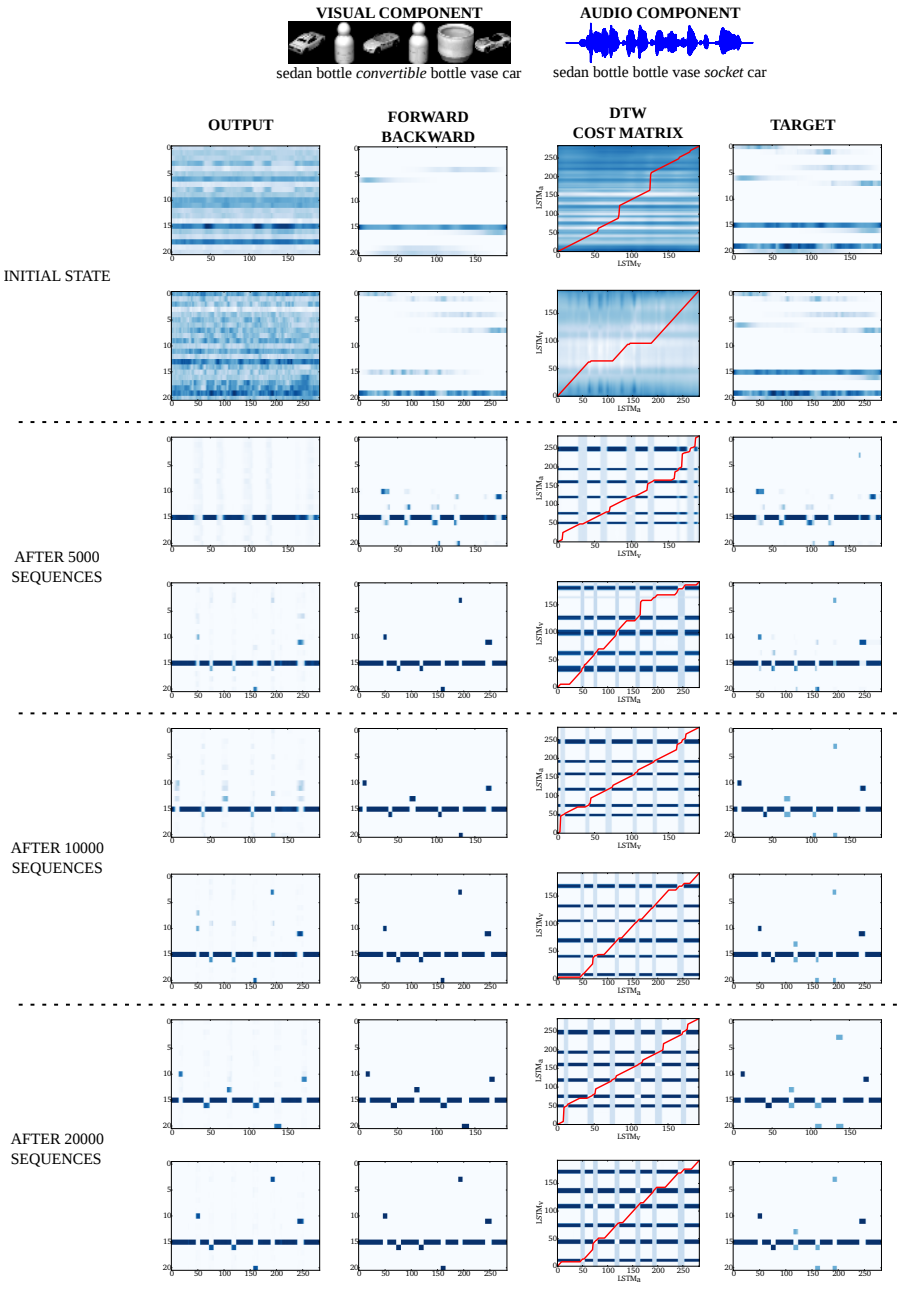

Figure 5: Example of several stages during the training. The top part of each row is the information from the visual component, whereas the bottom part is the information from the audio component. After training, both LSTMs agree on a similar symbolic feature for common semantic concepts, e.g. sedan, bottle.

concept "cup" (at the first row) is represented by the index "13" in both modalities and is classified at the first position of both sequences. In addition, not only the common elements, but also the missing elements are classified correctly.

Figure 5 illustrates the behavior of the training rule for multimodal sequences with missing elements. Initially, both LSTMs do not agree in a common relation. During training, both LSTMs are slowly converging to a common symbolic structure regardless of the non-continuity of the semantic sequence. Always, the blank class is the first component that is aligned (dark blue in outputs), in this

case after 5.000 sequences. Furthermore, regions appeared in the DTW cost matrix, and the semantic concepts have a clear representation (symbolic feature). For example, the concept "car" of audio LSTM (bottom part - 10.000 sequences) is represented by the index "10". After training (last row), both LSTMs successfully classify each input signal. The visual input ('sedan, bottle, *convertible*, bottle, vase, car') is represented by the indexes '10 16 13 16 20 11' (dark blue in output), whereas the audio input ('sedan bottle bottle vase *socket* car') is represented by the indexes '10 16 16 20 3 11'. These results indicate that both LSTMs agree on the same representation for the object-word association, even with missing elements ('convertible and socket') in the multimodal sequence.

7 CONCLUSIONS

In summary, we have presented a solution for the object-word association problem in a multimodal sequence with missing elements. Further work is planned for more realistic scenarios where the visual component is not clear segmented (noise in the background). Also, we are interested to extend the word-association problem between a two-dimensional image and speech. With this in mind, we will explore visual attention via eye-tracking in synchronization with speech. Finally, this task can be seen as simple, but it is still an open problem for understanding the language development in humans (Needham et al., 2005; Steels, 2008).

REFERENCES

Andersen, E. S., Dunlea, A., and Kekelis, L. The impact of input: language acquisition in the visually impaired. *First Language*, 13(37):23–49, January 1993. ISSN 0142-7237.

Asano, Michiko, Imai, Mutsumi, Kita, Sotaro, Kitajo, Keiichi, Okada, Hiroyuki, and Thierry, Guillaume. Sound symbolism scaffolds language development in preverbal infants. *cortex*, 63:196–205, 2015.

Berndt, Donald J and Clifford, James. Using Dynamic Time Warping to Find Patterns in Time Series. pp. 359–370, 1994.

Gershkoff-Stowe, Lisa and Smith, Linda B. Shape and the first hundred nouns. *Child development*, 75(4):1098–114, 2004. ISSN 0009-3920.

Graves, Alex, Fernández, Santiago, Gomez, Faustino, and Schmidhuber, Jürgen. Connectionist temporal classification. In *Proceedings of the 23rd international conference on Machine learning - ICML '06*, pp. 369–376, New York, New York, USA, 2006. ACM Press. ISBN 1595933832.

Harnad, Stevan. The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1):335–346, 1990.

Hochreiter, Sepp. The Vanishing Gradient Problem During Learning Recurrent Neural Nets and Problem Solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 06(02):107–116, April 1998. ISSN 0218-4885.

Hochreiter, Sepp and Schmidhuber, Juergen. Long Short-Term Memory. *Neural Computation*, 9 (8):1735—-1780, 1997.

Karpathy, Andrej, Joulin, Armand, and Li, Fei Fei F. Deep fragment embeddings for bidirectional image sentence mapping. In *Advances in Neural Information Processing Systems*, pp. 1889–1897, 2014.

Nakamura, Tomoaki, Araki, Takaya, Nagai, Takayuki, and Iwahashi, Naoto. Grounding of word meanings in latent dirichlet allocation-based multimodal concepts. *Advanced Robotics*, 25(17): 2189–2206, 2011.

Needham, Chris J, Santos, Paulo E, Magee, Derek R, Devin, Vincent, Hogg, David C, and Cohn, Anthony G. Protocols from perceptual observations. *Artificial Intelligence*, 167(1):103–136, 2005.

Nene, SA, Nayar, SK, and Murase, H. Columbia object image library (coil-20). Technical report, 1996.

Raue, Federico, Byeon, Wonmin, Breuel, T.M., and Liwicki, Marcus. Symbol Grounding in Multimodal Sequences using Recurrent Neural Network. In *Workshop Cognitive Computation: Integrating Neural and Symbolic Approaches at NIPS 05*, 2015.

Sohn, Kihyuk, Shang, Wenling, and Lee, Honglak. Improved multimodal deep learning with variation of information. In *Advances in Neural Information Processing Systems*, pp. 2141–2149, 2014.

Spencer, P E. Looking without listening: is audition a prerequisite for normal development of visual attention during infancy? *Journal of deaf studies and deaf education*, 5(4):291–302, January 2000. ISSN 1465-7325.

Srivastava, Nitish and Salakhutdinov, Ruslan R. Multimodal learning with deep boltzmann machines. In *Advances in neural information processing systems*, pp. 2222–2230, 2012.

Steels, Luc. The symbol grounding problem has been solved, so whats next ? *Symbols, Embodiment and Meaning. Oxford University Press, Oxford, UK*, (2005):223–244, 2008. ISSN 978-0199217274. doi: 10.1093/acprof:oso/9780199217274.003.0012.

Taylor, Paul, Black, Alan W, and Caley, Richard. The architecture of the festival speech synthesis system. 1998.

Vinyals, Oriol, Toshev, Alexander, Bengio, Samy, and Erhan, Dumitru. Show and tell: A neural image caption generator. *arXiv preprint arXiv:1411.4555*, 2014.

Yu, Chen and Ballard, Dana H. A multimodal learning interface for grounding spoken language in sensory perceptions. *ACM Transactions on Applied Perception (TAP)*, 1(1):57–80, 2004.